Deep Graph Learning: Foundations, Advances and Applications

Yu Rong*†
Tingyang Xu†
Junzhou Huang†
Wenbing Huang‡
Hong Cheng§
† Tencent AI Lab
‡ Tsinghua University
§ The Chinese University of Hong Kong
{royrong,tingyangxu,joehhuang}@tencent.com
huangwb@tsinghua.edu.cn
hcheng@se.cuhk.edu.hk

Yao Ma*¶
Yiqi Wang¶
Tyler Derr**
Lingfei Wu††
Tengfei Ma††
¶Michigan State University
** Vanderbilt University
††IBM Research AI
{mayao4,wangy206}@msu.edu
derrtylers@acm.org
{wuli,tengfei.ma1}@us.ibm.com

ABSTRACT

Many real data come in the form of non-grid objects, *i.e.* graphs, from social networks to molecules. Adaptation of deep learning from grid-alike data (*e.g.* images) to graphs has recently received unprecedented attention from both machine learning and data mining communities, leading to a new cross-domain field—Deep Graph Learning (DGL). Instead of painstaking feature engineering, DGL aims to learn informative representations of graphs in an end-to-end manner. It has exhibited remarkable success in various tasks, such as node/graph classification, link prediction, etc.

In this tutorial, we aim to provide a comprehensive introduction to deep graph learning. We first introduce the theoretical foundations on deep graph learning with a focus on describing various Graph Neural Network Models (GNNs). We then cover the key achievements of DGL in recent years. Specifically, we discuss the four topics: 1) training deep GNNs; 2) robustness of GNNs; 3) scalability of GNNs; and 4) self-supervised and unsupervised learning of GNNs. Finally, we will introduce the applications of DGL towards various domains, including but not limited to drug discovery, computer vision, medical image analysis, social network analysis, natural language processing and recommendation.

1 TARGET AUDIENCE AND PREREQUISITES

Target Audience: the conference attendees with interest in deep learning and graph mining.

Pre-requisites: for audiences who have the basic knowledge of deep learning and graph mining.

2 TUTORIAL TOPICS

As an emerging area with fruitful progresses in both theory and application recently, deep graph learning should be of interest to many of the conference attendees. The topics of this full-day tutorial include (but are not limited to) the following:

- Old-Fashioned Graph Learning [21]
- Foundations of Deep Graph Learning [7, 8, 14, 23]
- Advanced Deep Graph Learning [5, 17, 18]
  - Training deep GNNs [19, 22]
  - Scalability of GNNs [3, 10, 26]
  - Robustnes of GNNs [2, 4, 13, 28]
  - Self/Un-supervised learning of GNNs [9, 12, 20, 24]
- Applications: Drugs, Society, Vision, NLP and Recommendation [1, 6, 11, 15, 16, 25, 27]

3 TUTORS BIOGRAPHY

Yu Rong is a senior researcher of Machine Learning Center in Tencent AI Lab. He received his Ph.D. degree from The Chinese University of Hong Kong in 2016. He joined Tencent AI Lab in June 2017. His main research interests include graph neural networks and large-scale graph systems, with a particular focus on the design and efficient training for complex graph learning models. He has published several papers on data mining, machine learning top conferences, such as KDD, WWW, NeurIPS, ICLR, CVPR, ICCV, etc.

Wenbing Huang is now an assistant researcher at Tsinghua University. He received his Ph.D. degree of computer science and technology from Tsinghua University in 2017. His current research mainly lies in learning on irregular structures including graphs and videos. He has published about 30 top-tier conference and journal papers, including NeurIPS, ICLR, ICML, CVPR, etc.

Tingyang Xu is a senior researcher of Machine Learning Center in Tencent AI Lab. He obtained the Ph.D. degree from The University of Connecticut in 2017 and joined Tencent AI Lab in July 2017. In Tencent AI Lab, he is working on deep graph learning and applying the deep graph learning model to various applications, such as molecular generation and rumor detection. He has published several papers on data mining, machine learning top conferences KDD, WWW, NeurIPS, ICLR, CVPR, ICML, etc.

Hong Cheng is an Associate Professor in the Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong. She received the PhD degree from the University of Illinois at Urbana-Champaign in 2008. Her research inter-

*Co-first authors and were the leads of the two groups that equally contributed.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

ACM ISBN 978-1-4503-7998-4/20/08.
https://doi.org/10.1145/3394486.3406474

ests include data mining, database systems, and machine learning. She received research paper awards at ICDE'07, SIGKDD'06, and SIGKDD'05, and the certificate of recognition for the 2009 SIGKDD Doctoral Dissertation Award. She received the 2010 Vice-Chancellor's Exemplary Teaching Award at the CUHK.

Junzhou Huang is an Associate Professor in the Computer Science and Engineering department at the University of Texas at Arlington. He also served as the director of Machine Learning Center in Tencent AI Lab. His major research interests include machine learning, computer vision. He was selected as one of the 10 emerging leaders in multimedia and signal processing by the IBM T.J. Watson Research Center in 2010. His work won the MICCAI Young Scientist Award 2010, the FIMH Best Paper Award 2011, the MICCAI Young Scientist Award Finalist 2011, the STMI Best Paper Award 2012, the NIPS Best Reviewer Award 2013, the MICCAI Best Student Paper Award Finalist 2014 and the MICCAI Best Student Paper Award 2015. He received the NSF CAREER Award in 2016.

Yao Ma is a Ph.D. student of Computer Science and Engineering at Michigan State University. His research interests include network embedding and graph neural networks for graph representation learning. He was the leading presenter for the tutorial of "Graph Neural Networks: methods and applications" at AAAI2020 that is the most well-received tutorial with more than 400 audience. Updated information can be found at http://cse.msu.edu/~mayao4/.

Yiqi Wang is a Ph.D. student in the Computer Science and Engineering Department at Michigan State University. She is working on graph neural networks including fundamental algorithms, robustness and applications. She is one of the key contributors to the survey and empirical study on adversarial attacks and defenses on graphs with the developed repository. Updated information can be found at http://cse.msu.edu/~wangy206.

Tyler Derr is an Assistant Professor at Vanderbilt University in the Electrical Engineering and Computer Science department. He received his PhD in Computer Science from Michigan State University in 2020. His research is in network analysis and representation learning. He has published and serves as a program committee member at the top conferences in these domains and co-organized the Deep Graph Learning workshop at IEEE BigData'19. He received the Best Reviewer Award at ICWSM'19 and Best Student Poster Award at SDM'19. More details can be found at http://www.TylerDerr.com.

Lingfei Wu is a research staff member at IBM Research and is leading a research team (10+ RSMs) for developing novel Graph Neural Networks for many AL/ML/NLP tasks. He has published more than 60 top-ranked conference and journal papers and is a co-inventor of more than 25 filed US patents. He has co-organized 10+ conferences and workshops, including IEEE BigData'19, IEEE BigData'18, Workshops of Deep Learning on Graphs (with IJCAI'20, AAAI'20, KDD'19, and IEEE BigData'19). He has currently served as Associate Editor of ACM TKDD, and regularly served as a SPC/PC member of the following major AI/ML/NLP conferences including KDD, IJCAI, AAAI, NIPS, ICML, ICLR, and ACL. Personal website: https://sites.google.com/a/email.wm.edu/teddy-lfwu/home.

Tengfei Ma is a research staff member of IBM Research AI. Prior to moving to the IBM T. J. Watson Research Center in 2016, he received his Ph.D. from the University of Tokyo and joined IBM Research Tokyo in 2015. His research interests have spanned a number of different topics in machine learning and natural language processing. Recently his research is mainly focused on graph neural networks and their applications in healthcare and natural language processing; and he has published a series of papers about this topic in top conferences such as NeurIPS, ICLR, IJCAI, AAAI. Personal website: www.matengfei.com

REFERENCES

[1] Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor Detection on Social Media with Bi-Directional Graph Convolutional Networks. In *AAAI*.

[2] Heng Chang, Yu Rong, Tingyang Xu, Wenbing Huang, Honglei Zhang, Peng Cui, Wenwu Zhu, and Junzhou Huang. 2020. A Restricted Black-Box Adversarial Framework Towards Attacking Graph Embedding Models.. In *AAAI*. 3389–3396.

[3] Jie Chen, Tengfei Ma, and Cao Xiao. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. In *ICLR*.

[4] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial Attack on Graph Structured Data. *ICML* (2018), 1115–1124.

[5] Tyler Derr, Yao Ma, and Jiliang Tang. 2018. Signed graph convolutional networks. In *2018 IEEE International Conference on Data Mining (ICDM)*. 929–934.

[6] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *TheWebConf*.

[7] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *ICML*. 1263–1272.

[8] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *NeurIPS*. 1025–1035.

[9] Weihua Hu*, Bowen Liu*, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2020. Strategies for Pre-training Graph Neural Networks. In *ICLR*. https://openreview.net/forum?id=HJlWWJSFDH

[10] Wenbing Huang, Tong Zhang, Yu Rong, and Junzhou Huang. 2018. Adaptive sampling towards fast graph representation learning. In *NeurIPS*. 4558–4567.

[11] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2018. Junction Tree Variational Autoencoder for Molecular Graph Generation. In *ICML*. 2323–2332.

[12] Wei Jin, Tyler Derr, Haochen Liu, Yiqi Wang, Suhang Wang, Zitao Liu, and Jiliang Tang. 2020. Self-supervised Learning on Graphs: Deep Insights and New Direction. *arXiv preprint arXiv:2006.10141* (2020).

[13] Wei Jin, Yaxin Li, Han Xu, Yiqi Wang, and Jiliang Tang. 2020. Adversarial Attacks and Defenses on Graphs: A Review and Empirical Study. *arXiv preprint arXiv:2003.00653* (2020).

[14] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.

[15] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. 2018. Adaptive graph convolutional neural networks. In *AAAI*.

[16] Ruoyu Li, Jiawen Yao, Xinliang Zhu, Yeqing Li, and Junzhou Huang. 2018. Graph CNN for survival analysis on whole slide pathological images. In *MICCAI*.

[17] Yao Ma, Ziyi Guo, Zhaochun Ren, Eric Zhao, Jiliang Tang, and Dawei Yin. 2020. Streaming graph neural networks. *SIGIR* (2020).

[18] Yao Ma, Suhang Wang, Charu C Aggarwal, and Jiliang Tang. 2019. Graph convolutional networks with eigenpooling. In *KDD*. 723–731.

[19] Kenta Oono and Taiji Suzuki. 2020. Graph neural networks exponentially lose expressive power for node classification. *ICLR*.

[20] Zhen Peng, Wenbing Huang, Minnan Luo, Qinghua Zheng, Yu Rong, Tingyang Xu, and Junzhou Huang. 2020. Graph Representation Learning via Graphical Mutual Information Maximization. In *TheWebConf*. 259–270.

[21] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*. ACM.

[22] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2020. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *ICLR*.

[23] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.

[24] Petar Velickovic, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2019. Deep Graph Infomax. (2019).

[25] Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, Michael Witbrock, and Vadim Sheinin. 2018. Graph2seq: Graph to sequence learning with attention-based neural networks. *arXiv preprint arXiv:1804.00823* (2018).

[26] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931* (2019).

[27] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. 2019. Graph Convolutional Networks for Temporal Action Localization. In *ICCV*.

[28] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial Attacks on Neural Networks for Graph Data. In *KDD*. 2847–2856.